Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	1	("5557780").PN.	US-PGPUB; USPAT; USOCR	OR	OFF	2005/12/08 14:11
L2	1	1 and segment	US-PGPUB; USPAT	OR	ON	2005/12/08 14:28
L3	1	1 and header	US-PGPUB; USPAT	OR	ON	2005/12/08 14:30
L4	1	1 and "64"	US-PGPUB; USPAT	OR	ON	2005/12/08 14:32
L5	1	map adj3 metadata adj3 variable	US-PGPUB; USPAT	OR	ON	2005/12/08 14:38
L6	6	map near3 metadata near3 variable	US-PGPUB; USPAT	OR	ON	2005/12/08 14:39
L7	5	6 not 5	US-PGPUB; USPAT	OR	ON	2005/12/08 14:38
L8	7	map\$4 near3 metadata near3 variable	US-PGPUB; USPAT	OR	ON	2005/12/08 14:39
L9	1	8 not 6	US-PGPUB; USPAT	OR	ON	2005/12/08 14:39
L10	2384	(715/513).CCLS.	US-PGPUB; USPAT; USOCR	OR	OFF	2005/12/08 16:01
L11	4228	(709/223).CCLS.	US-PGPUB; USPAT; USOCR	OR	OFF	2005/12/08 16:01
L12	274	(715/515).CCLS.	US-PGPUB; USPAT; USOCR	OR	OFF	2005/12/08 16:01
L13	365	(715/523).CCLS.	US-PGPUB; USPAT; USOCR	OR	OFF	2005/12/08 16:01
L14	71	map\$4 with metadata with (virtual or variable)	US-PGPUB; USPAT	OR	ON	2005/12/08 16:02
L15	5	map\$4 with metadata with (virtual or variable) with source	US-PGPUB; USPAT	OR	ON	2005/12/08 16:02
L16	1	map\$4 with metadata with (virtual or variable) with target	US-PGPUB; USPAT	OR	ON	2005/12/08 16:02
L17	4	15 not 16	US-PGPUB; USPAT	OR	ON	2005/12/08 16:03
L19	389	map\$4 with source with virtual	US-PGPUB; USPAT	OR	ON	2005/12/08 16:03
L21	298	map\$4 with target with virtual	US-PGPUB; USPAT	OR	ON	2005/12/08 16:03

L22	50	19 and 21	US-PGPUB; USPAT	OR	ON	2005/12/08 16:03
L23	3	22 and metadata	US-PGPUB; USPAT	OR	ON	2005/12/08 16:04
L24	2	23 not 16	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L28	411	(electronic adj data adj interchange) with format	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L30	60	((electronic adj data adj interchange) with format) and metadata	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L31	2	(((electronic adj data adj interchange) with format) and metadata) and (virtual adj document)	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L32	17	(((electronic adj data adj interchange) with format) and metadata) and (map\$4 with document)	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L33	15	((((electronic adj data adj interchange) with format) and metadata) and (map\$4 with document)) not (((electronic adj data adj interchange) with format) and metadata) and (virtual adj document))	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L34	15	(((((electronic adj data adj interchange) with format) and metadata) and (map\$4 with document)) not ((((electronic adj data adj interchange) with format) and metadata) and (virtual adj document))) and xml	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L45	1	data adj transaction adj formatting adj standard	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L46	10	data with transaction with formatting with standard	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L47	9	(data with transaction with formatting with standard) not (data adj transaction adj formatting adj standard)	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L50	519	(edi) with (xml)	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L51	145	(edi) with (xml) with includ\$4	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L52	1	(edi) adj8 includ\$4 adj8 (xml)	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L53	78	(edi) with (includ\$4 adj8 (xml))	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05

L54	77	((edi) with (includ\$4 adj8 (xml))) not ((edi) adj8 includ\$4 adj8 (xml))	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L55	519	xml with edi	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L56	686	xml with advantage	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L57	30	(xml with edi) and (xml with advantage)	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L58	7	(self adj describ\$4 adj mark\$up adj language) with xml	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L59	1	((self adj describ\$4 adj mark\$up adj language) with xml) and (advantage with xml)	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L60	1	((self adj describ\$4 adj mark\$up adj language) with xml) and (benefit with xml)	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L61	6	((self adj describ\$4 adj mark\$up adj language) with xml) not (((self adj describ\$4 adj mark\$up adj language) with xml) and (advantage with xml)) not (((self adj describ\$4 adj mark\$up adj language) with xml) and (benefit with xml))	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L62	6	(((self adj describ\$4 adj mark\$up adj language) with xml) not (((self adj describ\$4 adj mark\$up adj language) with xml) and (advantage with xml)) not (((self adj describ\$4 adj mark\$up adj language) with xml) and (benefit with xml))) and advantage	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L63	4	(((self adj describ\$4 adj mark\$up adj language) with xml) not (((self adj describ\$4 adj mark\$up adj language) with xml) and (advantage with xml)) not (((self adj describ\$4 adj mark\$up adj language) with xml) and (benefit with xml))) and benefit	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L64	7	(self\$describ\$4 adj mark\$up adj language) with xml	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L65	73	(self with describ\$4 with mark\$up with language) with xml	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L66	66	((self with describ\$4 with mark\$up with language) with xml) not ((self\$describ\$4 adj mark\$up adj language) with xml)	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05

L67	12	(((self with describ\$4 with mark\$up with language) with xml) not ((self\$describ\$4 adj mark\$up adj language) with xml)) and (advantage with (language or xml))	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L68	2	edi with (self\$describ\$4 with mark\$up with language) with xml	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L78	1523	user with creat\$4 with map	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L79	11	map with metadata with variable	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L80	1	L78 and L79	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L81	442	map with based with name	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L82	37	L78 and L81	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L83	37	L82 not L80	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L85	7933	creat\$4 near2 map	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L86	61	L85 and L81	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L87	43	L86 not L83	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L88	43	L87 not L80	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L89	5	creat\$4 with map with (prior or before) with data with receiv\$4	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L90	1523	user with creat\$4 with map	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L92	3	user with creat\$4 with map with (prior or before) with receiv\$4	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L93	1	L76 and (stor\$4)	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L94	1	L76 and (database)	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L95	1	L76 and (predetermin\$4)	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L97	576	user with (creat\$4 near2 map)	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L98	518	creat\$4 with (prior or before) with data with receiv\$4	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L99	27	user with (creat\$4 near2 map) with first	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05

L100	137300	mapping	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L101	19	L99 and L100	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L102	281	map with based with (data or metadata) with (name or variable)	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L103	84	L102 and (creat\$4 with map)	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L104	23	L102 and (creat\$4 with map with user)	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L105	39	user with creat\$4 with map with receiv\$4 with (data or document)	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L106	19	L105 and mapping	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L107	18	L106 not L104	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L108	815	user with creat\$4 with mapping	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L110	305	user with (creat\$4 adj3 mapping)	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L111	453	map\$4 with based with (data or metadata or document) with (name or variable)	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L112	7	L110 and L111	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L113	795	map\$4 with based with name	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L114	22	L110 and L113	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L115	16	L114 not L112	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L116	4	user with (creat\$4 adj3 mapping) with (prior or before)	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L117	75	(map\$4 adj3 based) with (document or data or metadata) with (name or variable)	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L118	50	(map\$4 adj2 based) with (document or data or metadata) with (name or variable)	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L119	49	(map\$4 adj2 based) with (data or metadata) with (name or variable)	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L120	35	(map\$4 adj2 based) with (data or metadata) with (name)	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L121	4	database with predetermin\$4 with value with metadata	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05

L122	478	database with predetermin\$4 with value with data	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L123	189	database with (predetermin\$4 adj2 value) with data	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L124	189	L123 not L121	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L125	20	obtain\$4 with database with (predetermin\$4 adj2 value) with data	US-PGPUB; USPAT	OR	ON	2005/12/08 16:05
L129	34	10 and (map\$4 with metadata)	US-PGPUB; USPAT	OR	ON	2005/12/08 16:07



Subscribe (Full Service) Register (Limited Service, Free) Login

Search: • The ACM Digital Library O The Guide

map metadata variable virtual source target

### THE ACM DIGITAL LIBRARY

Feedback Report a problem Satisfaction survey

Terms used map metadata variable virtual source target

Found **64,558** of **167,655** 

Sort results by Display

results

relevance

expanded form

Save results to a Binder Search Tips Open results in a new

Try an Advanced Search Try this search in The ACM Guide

next

Results 1 - 20 of 200

window

Result page: 1 2 3 4 5 6 7 8 9 10

Relevance scale 🔲 📟 📟

Best 200 shown

Fast detection of communication patterns in distributed executions

Thomas Kunz, Michiel F. H. Seuren

November 1997 Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research

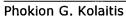
**Publisher: IBM Press** 

Full text available: pdf(4.21 MB)

Additional Information: full citation, abstract, references, index terms

Understanding distributed applications is a tedious and difficult task. Visualizations based on process-time diagrams are often used to obtain a better understanding of the execution of the application. The visualization tool we use is Poet, an event tracer developed at the University of Waterloo. However, these diagrams are often very complex and do not provide the user with the desired overview of the application. In our experience, such tools display repeated occurrences of non-trivial commun ...

Invited talk: Schema mappings, data exchange, and metadata management



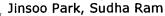
June 2005 Proceedings of the twenty-fourth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems

Publisher: ACM Press

Full text available: Plant (265.63 KB) Additional Information: full citation, abstract, references

Schema mappings are high-level specifications that describe the relationship between database schemas. Schema mappings are prominent in several different areas of database management, including database design, information integration, data exchange, metadata management, and peer-to-peer data management systems. Our main aim in this paper is to present an overview of recent advances in data exchange and metadata management, where the schema mappings are between relational schemas. In addition, w ...

Information systems interoperability: What lies beneath?



October 2004 ACM Transactions on Information Systems (TOIS), Volume 22 Issue 4

Publisher: ACM Press

Full text available: pdf(824.78 KB) Additional Information: full citation, abstract, references, index terms

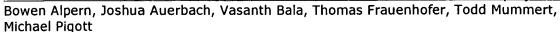
Interoperability is the most critical issue facing businesses that need to access information from multiple information systems. Our objective in this research is to develop a comprehensive framework and methodology to facilitate semantic interoperability among distributed and heterogeneous information systems. A comprehensive framework for



managing various semantic conflicts is proposed. Our proposed framework provides a unified view of the underlying representational and reasoning formalism ...

Keywords: Information integration, mediators, ontology, semantic conflict resolution, semantic heterogeneity

Distributed VEEs: PDS: a virtual execution environment for software deployment



June 2005 Proceedings of the 1st ACM/USENIX international conference on Virtual execution environments

Publisher: ACM Press

Full text available: pdf(299.26 KB) Additional Information: full citation, abstract, references, index terms

The Progressive Deployment System (PDS) is a virtual execution environment and infrastructure designed specifically for deploying software, or "assets", on demand while enabling management from a central location. PDS intercepts a select subset of system calls on the target machine to provide a partial virtualization at the operating system level. This enables an asset's install-time environment to be reproduced virtually while otherwise not isolating the asset from peer applications on the targ ...

Keywords: deployment, installation, management, streaming, virtualization

Process migration

September 2000 ACM Computing Surveys (CSUR), Volume 32 Issue 3

Publisher: ACM Press

Additional Information: full citation, abstract, references, citings, index Full text available: pdf(1.24 MB) terms, review

Process migration is the act of transferring a process between two machines. It enables dynamic load distribution, fault resilience, eased system administration, and data access locality. Despite these goals and ongoing research efforts, migration has not achieved widespread use. With the increasing deployment of distributed systems in general, and distributed operating systems in particular, process migration is again receiving more attention in both research and product development. As hi ...

Keywords: distributed operating systems, distributed systems, load distribution, process migration

Computing curricula 2001

html(2.78 KB)

September 2001 Journal on Educational Resources in Computing (JERIC)

Publisher: ACM Press

Full text available: pdf(613.63 KB) Additional Information: full citation, references, citings, index terms

Virtual machines: Scale and performance in the Denali isolation kernel

Andrew Whitaker, Marianne Shaw, Steven D. Gribble December 2002 ACM SIGOPS Operating Systems Review, Volume 36 Issue SI

Publisher: ACM Press

Full text available: pdf(1.91 MB) Additional Information: full citation, abstract, references, citings

This paper describes the Denali isolation kernel, an operating system architecture that





safely multiplexes a large number of untrusted Internet services on shared hardware. Denali's goal is to allow new Internet services to be "pushed" into third party infrastructure, relieving Internet service authors from the burden of acquiring and maintaining physical infrastructure. Our isolation kernel exposes a virtual machine abstraction, but unlike conventional virtual machine monitors, Denali does not ...

<sup>8</sup> Technical correspondence: Language integration in the common language runtime



Jennifer Hamilton

February 2003 ACM SIGPLAN Notices, Volume 38 Issue 2

Publisher: ACM Press

Full text available: P pdf(974.52 KB) Additional Information: full citation, abstract, references

The Common Language Runtime (CLR) is language and platform-neutral, and provides the underlying infrastructure for the Microsoft .NET Framework. A key innovation in the CLR is its support for multiple programming languages, enabling programming language integration at the runtime level to a much greater degree than is currently possible.

Keywords: common type system, exception handling, intermediate language, language interoperability, metadata, virtual machine

Model-driven development of Web applications: the AutoWeb system



Piero Fraternali, Paolo Paolini

October 2000 ACM Transactions on Information Systems (TOIS), Volume 18 Issue 4

Publisher: ACM Press

Full text available: pdf(6.94 MB)

Additional Information: full citation, abstract, references, citings, index

This paper describes a methodology for the development of WWW applications and a tool environment specifically tailored for the methodology. The methodology and the development environment are based upon models and techniques already used in the hypermedia, information systems, and software engineering fields, adapted and blended in an original mix. The foundation of the proposal is the conceptual design of WWW applications, using HDM-lite, a notation for the specification of structure, nav ...

Keywords: HTML, WWW, application, development, intranet, modeling

10 Industrial papers: metadata management for data integration: Clio grows up: from



research prototype to industrial tool

Laura M. Haas, Mauricio A. Hernández, Howard Ho, Lucian Popa, Mary Roth June 2005 Proceedings of the 2005 ACM SIGMOD international conference on Management of data

**Publisher: ACM Press** 

Full text available: pdf(339.65 KB) Additional Information: full citation, abstract, references

Clio, the IBM Research system for expressing declarative schema mappings, has progressed in the past few years from a research prototype into a technology that is behind some of IBM's mapping technology. Clio provides a declarative way of specifying schema mappings between either XML or relational schemas. Mappings are compiled into an abstract query graph representation that captures the transformation semantics of the mappings. The query graph can then be serialized into different query langua ...

11 Incommunicado: efficient communication for isolates



Krzysztof Palacz, Jan Vitek, Grzegorz Czajkowski, Laurent Daynas

November 2002 ACM SIGPLAN Notices, Proceedings of the 17th ACM SIGPLAN

#### conference on Object-oriented programming, systems, languages, and applications OOPSLA '02, Volume 37 Issue 11

Publisher: ACM Press

Full text available: 19 pdf(386.23 KB) Additional Information: full citation, abstract, references

Executing computations in a single instance of safe language virtual machine can improve performance and overall platform scalability. It also poses various challenges. One of them is providing a fast inter-application communication mechanism. In addition to being efficient, such a mechanism should not violate any functional and non-functional properties of its environment, and should also support enforcement of application-specific security policies. This paper explores the design and implement ...

**Keywords:** application isolation, inter-application communication

12 Shangri-La: achieving high performance from compiled network applications while





enabling ease of programming

Michael K. Chen, Xiao Feng Li, Ruiqi Lian, Jason H. Lin, Lixia Liu, Tao Liu, Roy Ju June 2005 ACM SIGPLAN Notices, Proceedings of the 2005 ACM SIGPLAN conference on Programming language design and implementation PLDI '05, Volume 40 Issue 6

**Publisher: ACM Press** 

Additional Information: full citation, abstract, references, citings, index Full text available: pdf(480.93 KB) terms

Programming network processors is challenging. To sustain high line rates, network processors have extremely tight memory access and instruction budgets. Achieving desired performance has traditionally required hand-coded assembly. Researchers have recently proposed high-level programming languages for packet processing, but the challenges of compiling these languages into code that is competitive with hand-tuned assembly remain unanswered. This paper describes the Shangri-La compiler, which acce ...

Keywords: chip multiprocessors, dataflow programming, network processors, packet processing, program partitioning, throughput-oriented computing

13 Secure virtual enclayes: Supporting coalition use of distributed application





technologies

May 2001 ACM Transactions on Information and System Security (TISSEC), Volume 4 Issue 2

**Publisher: ACM Press** 

Full text available: pdf(462.10 KB)

Additional Information: full citation, abstract, references, index terms, review

The Secure Virtual Enclaves (SVE) collaboration infrastructure allows multiple organizations to share their distributed application objects, while respecting organizational autonomy over local resources. The infrastructure is transparent to applications, which may be accessed via a web server, or may be based on Java or Microsoft's DCOM. The SVE infrastructure is implemented in middleware, with no modifications to COTS operating systems or network protocols. The system enables dynamic updates to ...

**Keywords**: Access control, coalition, collaborative system, group communication, middleware, security policy

14 Principled design of the modern Web architecture Roy T. Fielding, Richard N. Taylor





May 2002 ACM Transactions on Internet Technology (TOIT), Volume 2 Issue 2

**Publisher: ACM Press** 

Full text available: pdf(335.47 KB)

Additional Information: full citation, abstract, references, citings, index

The World Wide Web has succeeded in large part because its software architecture has been designed to meet the needs of an Internet-scale distributed hypermedia application. The modern Web architecture emphasizes scalability of component interactions, generality of interfaces, independent deployment of components, and intermediary components to reduce interaction latency, enforce security, and encapsulate legacy systems. In this article we introduce the Representational State Transfer (REST) arc ...

Keywords: Network-based applications, REST, World Wide Web

15 Concept-based querying in mediator systems

Kai-Uwe Sattler, Ingolf Geist, Eike Schallehn

March 2005 The VLDB Journal — The International Journal on Very Large Data Bases, Volume 14 Issue 1

Publisher: Springer-Verlag New York, Inc.

Full text available: pdf(329.34 KB) Additional Information: full citation, abstract

One approach to overcoming heterogeneity as a part of data integration in mediator systems is the use of metadata in the form of a vocabulary or ontology to represent domain knowledge explicitly. This requires including this meta level during query formulation and processing. In this paper, we address this problem in the context of a mediator that uses a concept-based integration model and an extension of the XQuery language called CQuery. This mediator has been developed as part of a project fo ...

**Keywords:** Data integration, Mediator systems, Query processing

16 Scalability, fidelity, and containment in the potemkin virtual honeyfarm



Michael Vrable, Justin Ma, Jay Chen, David Moore, Erik Vandekieft, Alex C. Snoeren, Geoffrey M. Voelker, Stefan Savage

October 2005 ACM SIGOPS Operating Systems Review, Proceedings of the twentieth ACM symposium on Operating systems principles SOSP '05, Volume 39 Issue

**Publisher: ACM Press** 

Full text available: pdf(506.39 KB) Additional Information: full citation, abstract, references, index terms

The rapid evolution of large-scale worms, viruses and bot-nets have made Internet malware a pressing concern. Such infections are at the root of modern scourges including DDoS extortion, on-line identity theft, SPAM, phishing, and piracy. However, the most widely used tools for gathering intelligence on new malware -- network honeypots -- have forced investigators to choose between monitoring activity at a large scale or capturing behavior with high fidelity. In this paper, we describe an approa ...

Keywords: copy-on-write, honeyfarm, honeypot, malware, virtual machine monitor

17 Research papers: data cleaning and mapping: Supporting executable mappings in



model management

Sergey Melnik, Philip A. Bernstein, Alon Halevy, Erhard Rahm

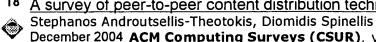
June 2005 Proceedings of the 2005 ACM SIGMOD international conference on Management of data

Publisher: ACM Press

Full text available: pdf(408.49 KB) Additional Information: full citation, abstract, references

Model management is an approach to simplify the programming of metadata-intensive applications. It offers developers powerful operators, such as Compose, Diff, and Merge, that are applied to models, such as database schemas or interface specifications, and to mappings between models. Prior model management solutions focused on a simple class of mappings that do not have executable semantics. Yet many metadata applications require that mappings be executable, expressed in SQL, XSLT, or other data ...

# 18 A survey of peer-to-peer content distribution technologies



December 2004 ACM Computing Surveys (CSUR), Volume 36 Issue 4

Publisher: ACM Press

Full text available: Phodf(517.77 KB) Additional Information: full citation, abstract, references, index terms

Distributed computer architectures labeled "peer-to-peer" are designed for the sharing of computer resources (content, storage, CPU cycles) by direct exchange, rather than requiring the intermediation or support of a centralized server or authority. Peer-to-peer architectures are characterized by their ability to adapt to failures and accommodate transient populations of nodes while maintaining acceptable connectivity and performance. Content distribution is an important peer-to-peer application ...

Keywords: Content distribution, DHT, DOLR, grid computing, p2p, peer-to-peer

## 19 Virtual machine monitors: Xen and the art of virtualization

Paul Barham, Boris Dragovic, Keir Fraser, Steven Hand, Tim Harris, Alex Ho, Rolf Neugebauer, Ian Pratt, Andrew Warfield

October 2003 Proceedings of the nineteenth ACM symposium on Operating systems principles

**Publisher: ACM Press** 

Additional Information: full citation, abstract, references, citings, index Full text available: pdf(168.76 KB) terms

Numerous systems have been designed which use virtualization to subdivide the ample resources of a modern computer. Some require specialized hardware, or cannot support commodity operating systems. Some target 100% binary compatibility at the expense of performance. Others sacrifice security or functionality for speed. Few offer resource isolation or performance quarantees; most provide only best-effort provisioning, risking denial of service. This paper presents Xen, an x86 virtual machine monit ...

Keywords: hypervisors, paravirtualization, virtual machine monitors

#### 20 An XML query engine for network-bound data

Zachary G. Ives, A. Y. Halevy, D. S. Weld

December 2002 The VLDB Journal — The International Journal on Very Large Data Bases, Volume 11 Issue 4

Publisher: Springer-Verlag New York, Inc.

Full text available: pdf(351.86 KB) Additional Information: full citation, abstract, citings, index terms

XML has become the lingua franca for data exchange and integration across administrative and enterprise boundaries. Nearly all data providers are adding XML import or export capabilities, and standard XML Schemas and DTDs are being promoted for all types of data sharing. The ubiquity of XML has removed one of the major obstacles to integrating data from widely disparate sources - namely, the heterogeneity of data formats. However, general-purpose integration of data across the wide are a also re  $\dots$ 







Keywords: Data integration, Data streams, Query processing, Web and databases, XML

Results 1 - 20 of 200

Result page: **1** <u>2</u> <u>3</u> <u>4</u> <u>5</u> <u>6</u> <u>7</u> <u>8</u> <u>9</u> <u>10</u>

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc. Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Mindows Media Player Real Player